

Characterization of Gigabit Ethernet Over Highly Turbulent Optical Wireless Links

*G.W. Johnson, J.P. Cornish, J.W. Wilburn, R.A. Young,
A.J. Ruggiero*

This article was submitted to
The International Society for Optical Engineering International
Symposium on Optical Science and Technology, Seattle,
Washington, July 7-11, 2002

U.S. Department of Energy

Lawrence
Livermore
National
Laboratory

July 1, 2002

DISCLAIMER

This document was prepared as an account of work sponsored by an agency of the United States Government. Neither the United States Government nor the University of California nor any of their employees, makes any warranty, express or implied, or assumes any legal liability or responsibility for the accuracy, completeness, or usefulness of any information, apparatus, product, or process disclosed, or represents that its use would not infringe privately owned rights. Reference herein to any specific commercial product, process, or service by trade name, trademark, manufacturer, or otherwise, does not necessarily constitute or imply its endorsement, recommendation, or favoring by the United States Government or the University of California. The views and opinions of authors expressed herein do not necessarily state or reflect those of the United States Government or the University of California, and shall not be used for advertising or product endorsement purposes.

This is a preprint of a paper intended for publication in a journal or proceedings. Since changes may be made before publication, this preprint is made available with the understanding that it will not be cited or reproduced without the permission of the author.

This report has been reproduced directly from the best available copy.

Available electronically at <http://www.doe.gov/bridge>

Available for a processing fee to U.S. Department of Energy
and its contractors in paper from
U.S. Department of Energy
Office of Scientific and Technical Information
P.O. Box 62
Oak Ridge, TN 37831-0062
Telephone: (865) 576-8401
Facsimile: (865) 576-5728
E-mail: reports@adonis.osti.gov

Available for the sale to the public from
U.S. Department of Commerce
National Technical Information Service
5285 Port Royal Road
Springfield, VA 22161
Telephone: (800) 553-6847
Facsimile: (703) 605-6900
E-mail: orders@ntis.fedworld.gov
Online ordering: <http://www.ntis.gov/ordering.htm>

OR

Lawrence Livermore National Laboratory
Technical Information Department's Digital Library
<http://www.llnl.gov/tid/Library.html>

Characterization of Gigabit Ethernet over highly turbulent optical wireless links

Gary W. Johnson, John P. Cornish, Jeffrey W. Wilburn,
Richard A. Young, and Anthony J. Ruggiero

Lawrence Livermore National Laboratory, P.O. Box 808, L-183,
Livermore, CA 94551

ABSTRACT

We report on the performance characterization and issues associated with using Gigabit Ethernet (GigE) over a highly turbulent ($C_n^2 > 10^{-12}$) 1.3 km air-optic lasercom links. Commercial GigE hardware is a cost-effective and scalable physical layer standard that can be applied to air-optic communications. We demonstrate a simple GigE hardware interface to a single-mode fiber-coupled, 1550 nm, WDM air-optic transceiver. TCP/IP serves as a robust and universal foundation protocol that has some tolerance of data loss due to atmospheric fading. Challenges include establishing and maintaining a connection with acceptable throughput under poor propagation conditions. The most useful link performance diagnostic is shown to be scintillation index, where a value of 0.2 is the maximum permissible for adequate GigE throughput. Maximum GigE throughput observed was 49.7% of that obtained with a fiber jumper when scintillation index is 0.1. Shortcomings in conventional measurements such as bit error rate are apparent. Prospects for forward error correction and other link enhancements will be discussed.

Keywords: Gigabit Ethernet, WDM, air-optic communications, free space optics, SATRN

1. INTRODUCTION

Government and military applications for optical wireless links call for deployments in extremely challenging environments. Establishing a link between ground, air, or sea-based mobile platforms in a wide variety of weather conditions leads to severe dynamic fading and pointing errors in addition to the usual propagation losses. To address these difficult situations, we began a new project called SATRN—Secure Air-optic Transport and Routing Network. Under the SATRN project we are currently developing a suite of new techniques including MEMS and electro-optical adaptive optics, optical signal processing, high-power fiber amplifiers, and forward error correction. Our intent is to improve link availability on long-range (>10 km) links by focusing on optical solutions, thus maintaining protocol transparency. Nearly all of these technologies intrinsically depend upon the link hardware being coupled in single-mode (SM) fiber, and that is the foundation for all of our transceiver optical hardware designs. We base our designs on SM fiber because it is highly scalable and transparent, and leverages commercial telecom components and standards.

To evaluate link performance under field conditions, we recently began experiments with Gigabit Ethernet (GigE) over our existing 1.3 km air-optic test range. GigE is a very popular, cost-effective communication standard that is already finding acceptance in commercial last-mile air-optic products. It is an excellent standard with which to evaluate link performance under conditions that include the high-level protocols required for real-world communications. We found it relatively easy to electrically connect it to the wavelength division multiplexed (WDM) optical channels that we currently use. A variety of protocols are available to suit the application. In particular, TCP/IP protocol tolerates most of the momentary connection failures typical of air-optic fading. The recent release of the 10 Gigabit Ethernet standard and a host of new products indicates a migration path that maintains the simplicity of interconnection. In the longer term, it appears that Ethernet may be destined to become the dominant standard in global networking¹, thus making its adoption in air-optic links imperative. The work reported here provides an initial benchmark from which the impact of technologies being developed in the SATRN program can be assessed.

2. SYSTEM OVERVIEW

Our primary test range is a 1.3 km round-trip, low-elevation, horizontal link located on-site at LLNL. All electro-optical components are optimized for ultradense WDM in the 1550 nm band and are coupled in SM fiber wherever possible. The 1550 nm band was chosen for its intrinsic eye safety and low propagation loss². This also allows us to apply standards and components from the telecommunications industry and simplifies the insertion of advanced electro-optical signal processing techniques. None of the advanced technologies under development were included in the system for this work, thus providing a worst-case baseline.

2.1 Horizontal 1.3 km Test Range

Figure 1 shows a view along the 1.3 km link looking from the transceiver station toward the remote fold mirrors. Beam height is 8 m above ground and several structures along with varying terrain can be seen, including a paved parking lot, grassy field, a building with boiler smokestacks, an evaporator tower, and a construction zone. Wind speeds are consistently high and tend to blow the boiler exhaust and evaporator mist across our line of sight. For reference, measurement of the refractive index structure parameter, C_n^2 , for this link has never dropped below 1×10^{-13} at any time over three years of observation, regardless of the measurement technique. It is consistently difficult to establish and maintain a reliable air-optic connection on this link at any time of day or night due to extreme scintillation and beam wander.




Figure 1 View along the beam line of the 1.3 km link. Note numerous structures and varying terrain on this 8-m altitude horizontal path.

The transmitter is a SM fiber-coupled, 20-cm, off-axis parabolic Newtonian telescope with a slow steering stage (0.05 Hz bandwidth) and visible camera for automatic alignment (Fig. 2). An image processing application keeps the transmitter pointed at the remote fold mirrors. The transmitted beam is 13.5 cm diameter (99% contour) and is focused midway to the receiver. Up to eight 2.5 Gbps directly-modulated laser transmitters are combined and fed to an erbium-doped fiber amplifier (EDFA) power amplifier operated at +28 dBm which then feeds the transmit telescope. We operated four wavelengths during these experiments, and all channels are power balanced within 1 dB through the entire system.

A fold mirror assembly is located 650 m from the transceiver station and consists of a pair of 25 cm diameter high reflectors, each mounted at 45 degrees and separated by 1 m. This provides a bi-directional path with uncorrelated

turbulence. One mirror has remotely-operated tip-tilt motors for rough alignment. The entire assembly is housed in a plywood box with weather-tight shutters and sits on a rooftop.

The receiver is based on a 30 cm diameter off-axis parabola with a 1.5 m focal length. Light from the primary first hits a slow steering mirror (0.05 Hz bandwidth) that is controlled in a manner similar to that of the transmitter, with the visible image derived from a video camera later in the receiver package. Next in line is a fast steering mirror with a 500 Hz bandwidth (Ball Aerospace model TT-25) that compensates for beam wander. Feedback is derived from a Ge position-sensitive detector and amplifier (Ontrak Photonics). At least 80% of the received 1550 nm energy is focused on a 63 μm multimode (MM) or 8 μm SM receive fiber.




Figure 2 Simplified plan view of transmitter and receiver optics, approximately to scale.

Received light is first amplified by a gain-flattened EDFA preamp (MPB Communications model EFA-R35W) with 36 dB of gain and a 4 dB noise figure. Its output then passes through a proprietary photonic-based level controller. A WDM filter separates the wavelengths for use by the 2.5 Gbps avalanche photodiode receivers. Threshold sensitivity for the receivers is -42 dBm (63 nW) referred to the SM input fiber. This translates to 394 photons per bit at 1.25 Gbps, the GigE speed. Clock and data recovery (CDR) is performed by Silicon Laboratories Si5023 chips. They include integral limiting amplifiers and accommodate all industry-standard data rates between 155 Mbps to 2.5 Gbps. This gives us great flexibility, including operation at heterogeneous rates.

2.2 Gigabit Ethernet Interface

An electrical interface to Gigabit Ethernet was fashioned from a GBIC evaluation board (Finisar FDB-1012) with limiting amplifiers (Maxim MAX3265) to fix both the transmit and receive levels, and broadband amplifiers (Picosecond Pulse Labs 5827) to supply the appropriate drive level to the laser transmitters. Eye diagrams were clean with a transmitted jitter of <22 ps rms. Two Dell 530 workstations were connected to the private GigE Ethernet. The workstations included dual 1.4 GHz processors and SCSI 160 RAID-0 disk arrays for maximum throughput, and ran Microsoft Windows 2000 SP2 with its standard TCP stack. Network interfaces were Syskonnect SK-9844 1000Base-SX dual-link cards. These support a backup link that automatically activates when the primary link fails. We did not use this feature in our tests, though it offers an interesting path toward implementing a high-reliability hybrid RF/optical communications link. Traffic was generated by running a simple TCP/IP file transfer application written in LabVIEW. Throughput over a direct fiber link between workstations when running that application was a steady 96 Mbps, which is 9.6% of theoretical. No attempt was made to optimize throughput.

2.3 Diagnostics

All diagnostic data was recorded simultaneously on all systems with absolute time synchronization within 1 s to allow correlation analysis. Bit error rate (BER) testing was performed with a Bitalzyer model BA-3600 (SyntheSys Research) which provides direct measurement of error burst length, error-free interval, and Q factor analysis. It also stores error data and permits re-analysis with optional application of forward error correction (FEC) codes to evaluate their performance. Eye diagrams were displayed on a Tektronix CSA-8000, and optical spectra and power levels were

observed on an Agilent 86140B. Analog values for received power and position-sensitive detector position data were acquired with a 16-bit data acquisition board in a PC running at 1000 samples per second, streaming data to disk continuously. Received power measurements were later analyzed in the following manner. Fade rate and fade fraction were measured by counting 1 ms samples while the power dipped below the known photoreceiver threshold. Fade fraction is the summed fade duration over a 1 s observation interval. Fade rate is defined as the frequency in Hz at which fades (of any duration) occur during the selected observation interval.

An effective in-fiber scintillation index was computed from the log variance of received intensity:

$$\sigma_{\ln I} = \frac{1}{4} \ln \left(\frac{\sigma_I^2}{\langle \bar{I} \rangle^2} \right) \quad (1)$$

Where $\langle \bar{I} \rangle$ is the mean intensity, and σ_I^2 is the variance of the intensity. An effective in-fiber refractive index structure parameter is computed from⁶

$$C_n^2 = \frac{C}{4} \frac{\sigma_I^2}{\langle \bar{I} \rangle^2} D^{7/3} L^{-3} \quad (2)$$

Where D is the transmitter aperture in meters, L is the path in meters, and C is the ratio of the receiver aperture to the transmitter aperture. For confirmation, C_n^2 was also computed from meteorological observations on a tower located on the LLNL site. Poor correlation between path-averaged and regional measurements are expected because of strong local perturbations along the optical path.

Ethernet performance was measured with EtherPeek, a Windows-based application from WildPackets, Inc., that reports throughput, packet size, packet error rate, and many other statistics. We also used a Finisar Gigabit Traffic Check as an excellent visual diagnostic for the GigE network. It displays performance and error status in both directions and uses standard GBICs for both network connections.




Figure 3 Received eye diagram with an active GigE connection under the best observed conditions. Upper trace, raw signal from the APD photoreceiver. Lower trace, output from the CDR circuit.

3. EXPERIMENTAL RESULTS

For all experiments, the system was configured with two WDM channels for GigE and two channels for BER testing at 1.25 Gbps. Total transmitter power was +28 dBm. Data was collected over a total of seven days during clear, sunny weather and on clear, cold nights during April and May, 2002. Figure 3 shows received eye diagrams under the best link conditions. Received jitter on the raw photoreceiver signal was 210 ps p-p, reduced to 50 ps after the clock/data recovery circuit. This easily meets the 1000BASE-CX receiver specification of 528 ps.

Received signal intensity variations were severe, even under the best conditions near midnight. In Fig. 4, a 10 s record of total received power in the SM fiber is plotted for two conditions. In the left graph, taken at about 5 PM, turbulence was severe, scintillation index was 0.26, and the GigE link could barely attain a connection for even a fraction of a second. Note the very large dynamic range, exceeding 40 dB. In contrast, the right graph, taken just before midnight, had a much better scintillation index of 0.1 and reliable GigE data transfer occurred. At that time, received power was sufficient to frequently saturate this particular power measurement receiver.




Figure 4 Total received power at two times of day. Left, in high turbulence at 5 PM, and right, minimum turbulence near midnight. No GigE conversation was possible in the high turbulence case. Sample rate is 1 kHz.

Results of analysis for fade rate and fade fraction on the received power analog signals are shown in Fig. 5. Both data sets encompassed 1.5 hours of 1 kHz analog data. For a perfect connection, fade rate would be a spike at zero Hz and fade fraction would similarly be zero. Under turbulent conditions (left graphs), fades occur at up to 100 Hz with a large number of fades as long as 100 ms. In all data sets, fades nearly 1 s long are observed. Atmospheric turbulence causes beam broadening, beam steering, and speckle. Beam broadening reduces the average incident intensity. Beam steering and scintillation (due to speckle) cause signal fades or dropouts. Long and frequent fades prove a need for highly robust forward error correction, data redundancy, or data retransmission if high link availability is required.

To further compare link performance, data was analyzed from five separate times of day and is summarized in Table 1. GigE throughput was compared to an ideal fiber link and we attained 47.9% of ideal under the best conditions. Even then, packet error rates and throughput varied wildly over time as one might expect. Also, the Windows networking software would report "network cable unplugged" whenever a sufficiently long fade occurred. Recovery time for the network was typically less than 1 s and data transfer would promptly resume. Since there is no error correction in standard internet protocol (IP), even a single bit error results in a packet error and retransmission is requested. For the hours of data where a usable GigE throughput was obtained, we measured an average 10.6% packet error rate. Contrast that with commercial fiber networks where packet error rates are vanishingly small, on the order of 10^{-9} or less.




Figure 5 Fade rate and fade fraction for two times of day. Left, in high turbulence at 5 PM, and right, minimum turbulence near midnight.

Table 1 Comparative Link Statistics

Run date & time	GigE Throughput	Avg. BER	Received Power(a)	Link Margin	Avg. Fade Fraction	Scint. Index	Fiber Cn ²	Met Tower Cn ² (b)
5-13 11:30	No conn.	2.0E-3	7 μ W	20.5 dB	0.039	0.52	2.8E-11	4E-13
5-14 13:45	Barely conn.	2.5E-4	8 μ W	21.0 dB	0.013	0.28	5.9E-12	3E-13
5-17 16:45	Barely conn.	4.0E-4	12 μ W	22.8 dB	0.0014	0.26	4.7E-12	<3E-14
5-17 20:51	19.4%	4.3E-6	35 μ W	27.4 dB	<.001	0.15	2.3E-12	N/A
5-17 23:10	47.9%	6.2E-6	40 μ W	28.0 dB	0.0065	0.12	1.9E-12	N/A

3 Average power per wavelength in single-mode receiver fiber

4 Met tower located about 500 m from the beam path, in an open field

The reason that actual throughput was reduced by almost 50% in the presence of only a 10.6% packet error rate is twofold. First, there were an unmeasured number of disconnections that required significant time for automatic recovery. Second, TCP uses *sliding windows* of variable size to manage acknowledgement messages along the reverse channel³. This allows streams of packets to saturate a reliable, uncongested network without requiring an acknowledgement after every single packet. If the network becomes unreliable, the window closes down and the handshaking process starts anew, gradually accelerating as the window lengthens in response to an effective data flow. The situation is even worse when the bandwidth*delay product is high, as it is in fast, long-range, air-optic links. Multiple packets can be in flight--and possibly lost--simultaneously. Newer additions to the IEEE 802.3 standard, such as Selected Acknowledgement (SACK), attempt to improve performance on long, fast links by improving retransmission behavior. This is not a problem for full packets on our 1.3 km link, since the propagation time is only 4.3 μ s.

One of our objectives is to determine link quality statistics that relate well with GigE performance. Average BER is a possible candidate, and data shows that 1×10^{-5} is an approximate upper limit for a usable network connection. However, it has already been established that classic BER data is of limited value here because the link performance is dominated by long error bursts including loss of synchronization in the BER instrument.⁴

Average received power is a possible diagnostic because it increases as scintillation decreases, all other things being equal. For the present data, adequate GigE performance required an average power at least 27 dB above the photoreceiver threshold. This link margin accounts for most of the fade depth already noted.

Perhaps the best performance estimator is scintillation index, which is derived from the ratio of variance to mean of received power as shown in Eqn. 1. This statistic is simple to measure and trends nicely with other quantitative and qualitative measures of link performance. Our data shows that scintillation index must be less than 0.2 in order to obtain adequate GigE performance. We have also made this measurement in real-time with a simple data acquisition system.

Our in-fiber measure of C_n^2 is related to scintillation index but has the added advantage of permitting comparison to other path-averaged C_n^2 measurements.⁵ Correlation with regional C_n^2 measurements, such as our local meteorological tower, is typically poor because of strong dependence on localized turbulence. Conditions near the met tower are benign in contrast to our beam path, leading to a much lower C_n^2 in the open field. Since our beams are coherent, and we are operating in such high scintillation regime, this measurement is probably near saturation most of the time.⁶

3.1. Performance of Forward Error Correction

Forward error correction (FEC) is often suggested as a means toward improving error rate for any communication system, and it is worth investigating the performance of FECs on our air-optic link data. Using the Bitalzyer, we analyzed stored data for a 4 minute period when the connection was running at its best, around midnight. Burst error analysis showed that the uncorrected average BER was 3.25×10^{-6} and that only 0.06% of the errors were single-bit errors; the remainder were bursts of two or more bits. We then applied a standard Reed-Solomon (255,223) FEC to the same data. This code can correct up to 16 bit errors in a 255-byte block. BER was reduced 12% to 2.89×10^{-6} . Considering the 14% overhead of the RS coding, we actually have an effective coding loss, rather than a gain.

Block error performance histograms show the effectiveness of the chosen FEC. In Fig. 6, the block size is set to 2040 bits to match our RS block size. Skew in the histograms indicates that the majority of blocks had more than 16 error bits within a 2040-bit block. Hence, this code is relatively ineffective for the present situation. While there are certainly more effective FECs available, they come at the expense of a great computational burden. At very high data rates such codes quickly become impractical to implement. Nevertheless, simple FECs offer a necessary and useful improvement for links with low BERs when fade conditions are less severe.




Figure 6 Block error histograms for a four-minute data set without (left) and with (right) the application of a Reed-Solomon (255,223) FEC.

4. CONCLUSION

Observations on our uncorrected, highly-turbulent 1.3 km link show that a GigE connection can only be maintained under the best atmospheric conditions. Analog measurement statistics on power in the receiver fiber can provide a useful real-time estimate of overall link performance, and can be used to estimate the likelihood of success in maintaining a GigE connection. For our receiver, an adequate link exhibits an in-fiber scintillation index less than about 0.2 and average power at least 27 dB above the photoreceiver threshold. This allows adequate link margin for fades. An average BER less than 1×10^{-5} is also an observable threshold, though it can only be obtained over long periods of time.

Planned insertion of our new optical technology is expected to provide orders of magnitude improvement in link availability. Our goal is to sufficiently improve link performance at the physical (electro-optical) layer and maintain transparency while avoiding solutions that require development of custom GigE hardware or software. This work is currently underway and will be reported on in later publications. It may be possible to improve Ethernet performance by adding custom interface hardware with complex forward error correction and/or added data redundancy. This would require a significant engineering effort and, depending upon the type and amount of data redundancy, could be an expensive and bulky subsystem. We also suspect that the TCP/IP protocol stack at the lower levels could be modified to better handle frequent connection loss. This might be implemented as a network processor, similar to a router. But those solutions are protocol-dependent and possibly not scalable.

We have begun measurements on a 29 km (each way), 2-degree angle slant path link between LLNL and Mt. Diablo. In these tests, a mobile transceiver system with a dome, telescopes, tracking mount, and support van were deployed and have established a 10 Gbps WDM link with the LLNL transceiver. We are planning on GigE experiments over this link in the near future. The mobile system will, in the future, be used for link experiments with mobile platforms such as unmanned aerial vehicles to demonstrate the effectiveness of wideband, secure, air-optic links in Government and military applications.

5. ACKNOWLEDGEMENTS

The authors would like to thank the many support personnel who designed, built, and operated the electro-optical equipment, especially Diane Cooke, Jeff Cooke, Noemi Fortes, Steve Mostek, and Dean Rippee. Important advice was also received from John Henderson. This work was performed under the auspices of the U.S. Department of Energy by Lawrence Livermore National Laboratory under Contract W-7405-ENG-48.

6. REFERENCES

1. J. Thatcher, "How optical Ethernet is disrupting the network marketplace: Understanding the causes, trends, and implications." *Proceedings of OFC 2002*, paper ThT (2002).
2. D. A. Rockwell and G. S. Mecherle, "Wavelength selection for optical wireless communications systems." *Proceedings of SPIE Vol. 4530*, Denver, August, 2002, pp. 27-35.
3. D. E. Comer, *Internetworking with TCP/IP*, 4th Edition, Prentice-Hall, Upper Saddle River, NJ, 1995.
4. P. Polak-Dingels, et. al., "Performance update of a terrestrial laser communication link with quadrature amplitude modulation." *Proceedings of SPIE Vol. 2990*, 1997, pp. 103-113.
5. S. C. Wilks, et. al., "Modeling of adaptive optics-based free-space communications systems." *Proceedings of SPIE Vol. 4821*, Seattle, July, 2002, paper no. 4821-15.
6. T. Wang, G. R. Ochs, and S. F. Clifford, "A saturation-resistant optical scintillometer to measure C_n^2 ." *J. Opt. Soc. Amer.*, Vol. 683(3), March, 1978, pp., 334-338.